AAS 2000-3939

EVENT DETECTION AND IDENTIFICATION DURING AUTONOMOUS INTERPLANETARY NAVIGATION

Timothy Crain*, Robert H. Bishop †, and Todd A. Ely ‡

An autonomous/adaptive interplanetary navigation architecture employing neural networks and genetic algorithms is being developed that will alert navigators when a shift from optimal to suboptimal filtering occurs and assist in modifying the filter parameters to resume optimal tracking. This architecture consists of a near real-time autonomous monitoring component and an offline adaptive component. The autonomous component analyzes operational filter residuals to detect the transition to suboptimal filtering and to identify the nature of the mismodeling. Once identified, the adaptive component then modifies the necessary model parameters to bring the filter back into optimal operation. The autonomous identification of mismodeling employs a hierarchical mixture-of-experts model where the experts are extended Kalman filters. The filters in the hierarchy are organized into banks and regulated by two levels of single layer neural networks called gating networks. The architecture of the overall navigation approach is introduced and the operation of the autonomous monitoring component is demonstrated. Two experiments will show the autonomous navigation component can successfully identify discrete model changes such as impulsive maneuvers and continuous model changes such as solar radiation mismodeling. A third experiment will demonstrate the robust decision making capability of the hierarchical mixture-of-experts by successfully identifying three successive model changes. All experiments are performed on Mars Pathfinder two way Doppler data for the period from February 4, 1997 to April 17, 1997.

*Graduate Student, The Center for Space Research, The University of Texas at Austin, TX 78759-5321, Member AAS, Member AIAA.

†Associate Professor, The Center for Space Research, The University of Texas at Austin, TX 78759-5321. Member AAS, Associate Fellow AIAA.

‡Staff Engineer, Navigation and Mission Design Section, Jet Propulsion Laboratory, Mail Stop 301-125L, 4800 Oak Grove Drive, Pasadena, CA 91109-8099.

INTRODUCTION

Current tracking of interplanetary spacecraft relies heavily upon navigator experience and ad hoc techniques for the resolution of anomalous residual signatures produced by the operational tracking filter. These anomalies are the result of changes in the dynamics and measurement system that cause the spacecraft environment to deviate from the operational tracking model. Difficulties navigating recent missions, such as modeling the solar radiation pressure (SRP) surface on Mars Pathfinder (MPF) [1] and the performance of Mars Climate Orbiter, illustrate the need for a systematic method to resolve anomalous behavior in the tracking solution. An autonomous/adaptive navigation architecture employing neural networks and genetic algorithms is being developed that will alert the navigator when a change from optimal to suboptimal filtering occurs and assist in modifying the filter parameters to resume optimal tracking. This architecture consists of a near real-time autonomous monitoring component and an offline adaptive component, as illustrated in Figure 1. The autonomous monitoring component operates in parallel with the operational filter, processing measurements as they become available. Operational filter residuals are compared to the residuals from a set of hypothetical filters contained in a hierarchical mixture-of-experts (HME) framework [2, 3]. If a transition to suboptimal modeling is detected, the nature of the mismodeling responsible for the transition is identified. Once identified, the adaptive component then modifies the necessary model parameters in an offline process to bring the operational filter back into optimal operation and the autonomous monitor is reset to process new data. Chaer and Bishop [4] proposed a genetic algorithm filter adaptation that used the HME as the performance index for parameter modification. However, this approach was found to evaluate performance relative only to filters in the HME and Ely, Bishop, and Crain [5] reformulate the adaptive component to use an absolute performance index based upon a sample statistic of filter residual histories.

All filters used in this work are extended Kalman filters (EKFs) processing MPF two-way Doppler DSN data. The dynamic and measurement models used to define each filter in the HME are variations of the operational model developed by the MPF navigation team detailed in the post-mission navigation report [1]. The current document is organized by a discussion of the architecture and theory of the HME, a discussion of experimental results, and a brief conclusion. Experiments were performed to test the autonomous monitor's ability to detect and identify: (1) small unmodeled impulsive maneuvers, (2) changes in the SRP model and impulsive maneuvers, and (3) unmodeled impulses and noise increases occurring in the same data span.

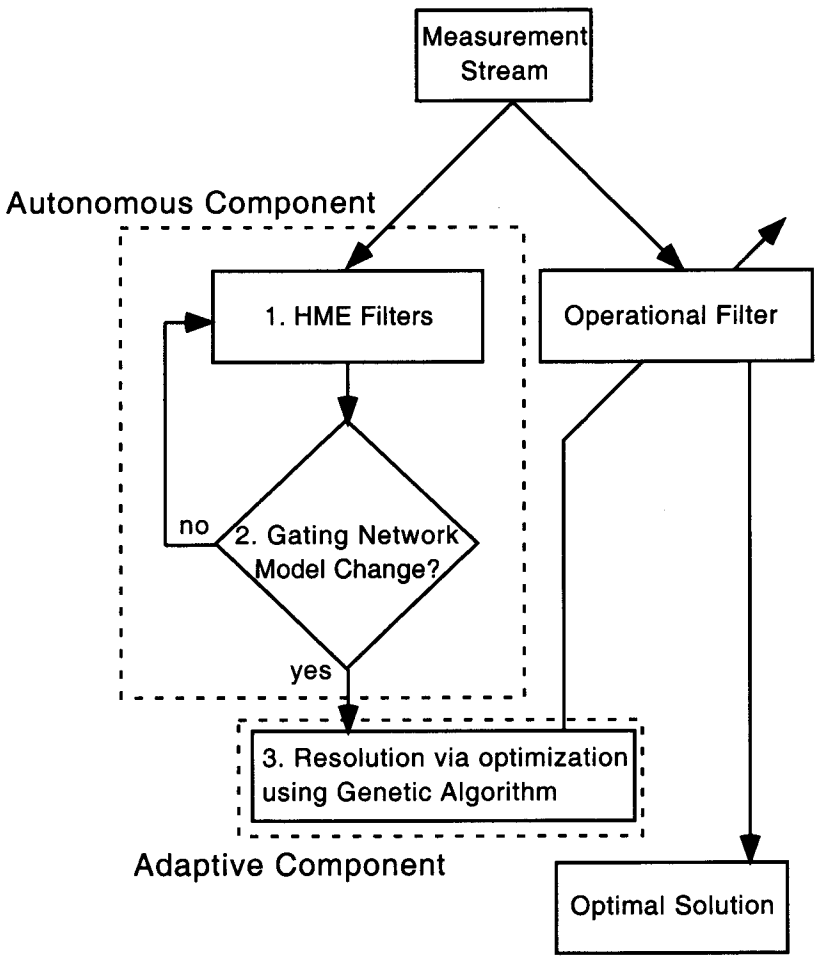

Figure 1: Autonomous/Adaptive Navigation Architecture

AUTONOMOUS MONITORING ARCHITECTURE

Each filter in the HME contains a unique measurement and dynamics model which can be represented by the parameter vector $\boldsymbol{\alpha}$. Thus, a particular filter realization may be thought of as an expert tracking system for a specific region of the modeling parameter space. The HME model is composed of macromode identification banks on the top level which are in turn composed of micromode identification experts (filters). Macromode environment changes include unmodeled impulsive maneuvers, SRP model changes, and measurement system noise changes. Filter realizations representing macromode changes in the spacecraft environment are collected into competing banks on the top level. Figure 2 illustrates the general HME configuration used in the experiments presented this study. The HME is separated into four filter banks. Three of the banks represent macromode changes in the spacecraft environment and contain multiple filters which model specific impulsive events, SRP model changes, and noise level changes. The fourth bank contains only the operational filter as an experimental control.

The HME is regulated by two levels of single layer neural networks known as gating networks (GN) as indicated in Figure 2. The bank level GNs assign weights to filters within a single bank to indicate the relative residual performance of each realization. The top level GN assigns weights to each bank according to the collective residual performance of the filters within the bank. Optimal tracking is assumed as long as the maximum value of the top level gating vector $\boldsymbol{g}_k$ at the time, t_k, of a measurement z_k is associated with the control bank (i.e. $g_{3,k}$ is the maximum element of $\boldsymbol{g}_k$). A shift to suboptimal tracking is indicated when the maximum element of $\boldsymbol{g}_k$ is associated with one of the banks representing an environment change macromode. Such a shift could be the results of a discrete event, as in the case of a thruster misfire, or a continuous or sustained environment change, as in the case of a steady change in shading of surfaces exposed to SRP. From a practical standpoint, a thruster misfire might only require a few hours worth of tracking to identify it, while a subtle change in SRP effects might take several days or weeks of measurement data to identify the change. Identification is accomplished by inspection of the macromode consistently given the highest top level gating weight. Once the environment change macromode has been identified, the adaptive component then operates in an offline mode to adjust the parameter set via a genetic algorithm. Finally, the operational filter is updated and the autonomous monitoring process resumes with the adapted optimal operational filter [5].

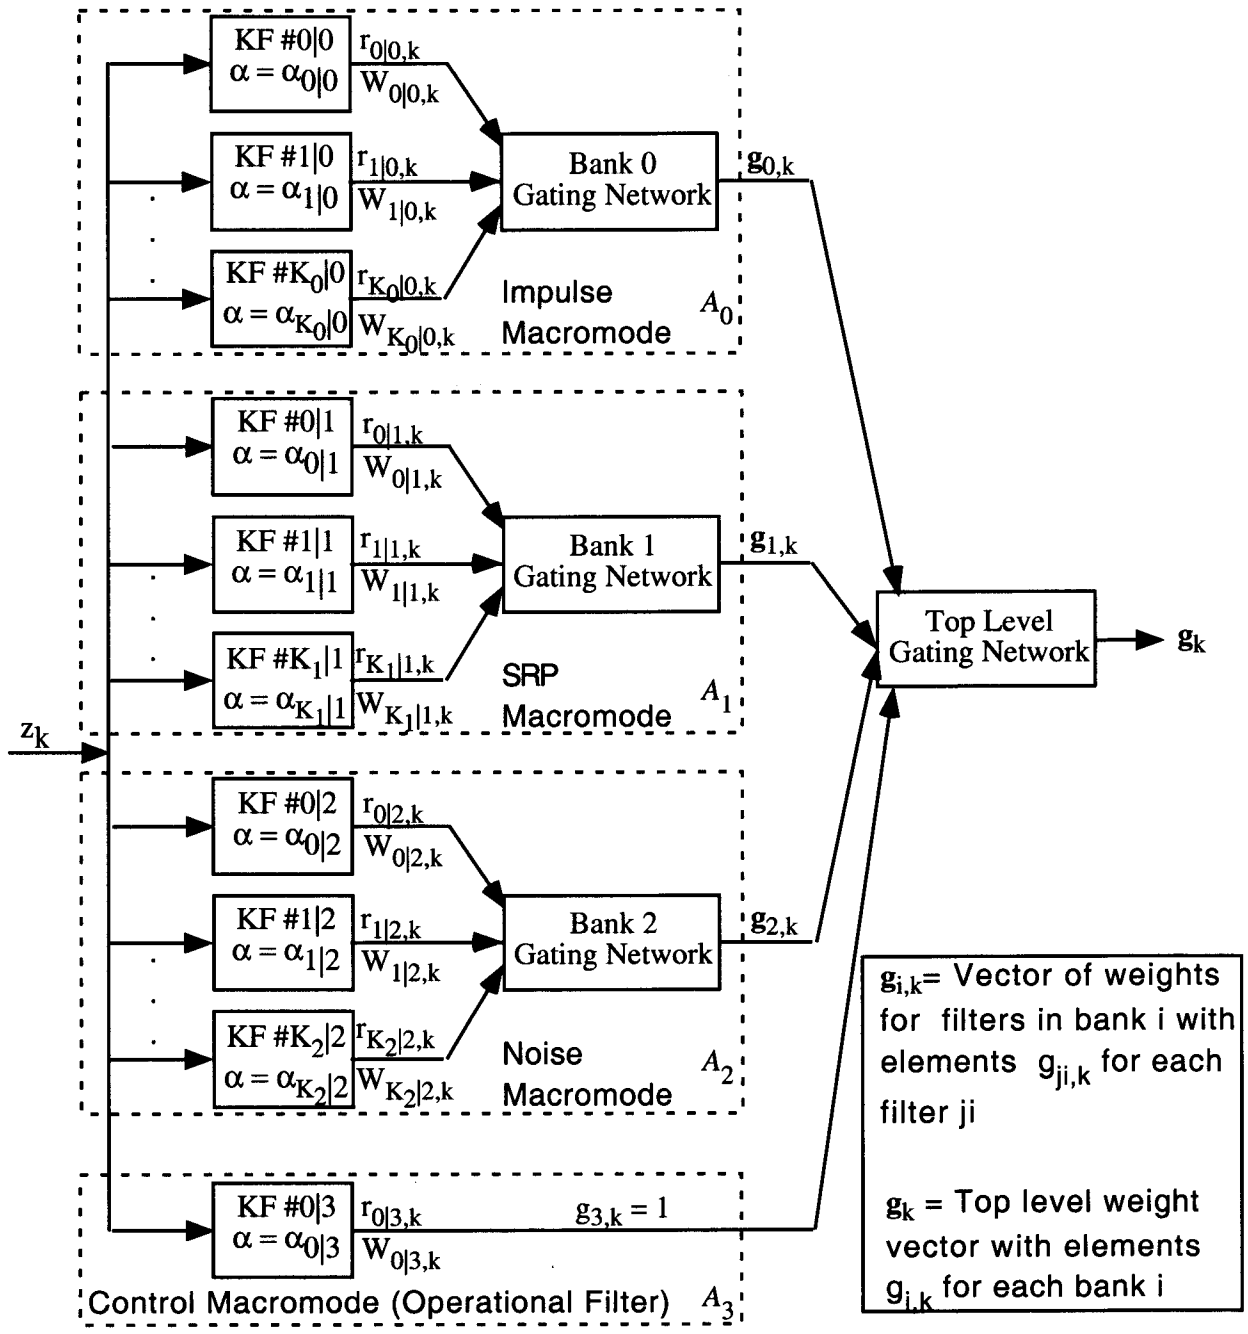

Figure 2: HME Configuration

THE HIERARCHICAL MIXTURE OF EXPERTS

The operation of the HME model regulated by multilevel GNs is based upon maximizing the probability density of the current input measurement, z_k, as approximated by a weighted combination of the conditional probabilities of each of the L macromodes in the HME

$$\begin{aligned} f(z_k \mid Z_{k-1}) &= \sum_{i=1}^{L} f(z_k \mid A_i, Z_{k-1}) P(A_i \mid Z_{k-1}) \\ &= \sum_{i=1}^{L} f(z_k \mid A_i, Z_{k-1}) g_{i,k} \end{aligned} \tag{1}$$

where A_i is the set of all filter realizations in the i^{th} macromode,

$$A_i = \{\boldsymbol{\alpha}_{ji} \mid 0 < j \leq K_i, 0 < i \leq L\}, \tag{2}$$

K_i is the number of filters in the i^{th} bank, and Z_{k-1} represents the set of all previous measurements. The conditional probability of the i^{th} macromode is $f(z_k \mid A_i, Z_{k-1})$ and $P(A_i \mid Z_{k-1})$ is the *a priori* probability that the i^{th} macromode models the input generating environment. The top level GN approximates the *a priori* probability using the gating weight $g_{i,k}$.

The conditional probability of the i^{th} macromode is given by

$$\begin{aligned} f(z_k \mid A_i, Z_{k-1}) &= \sum_{j=1}^{K_i} f(z_k \mid \boldsymbol{\alpha}_{ji}, Z_{k-1}) P(\boldsymbol{\alpha}_{ji} \mid Z_{k-1}) \\ &= \sum_{j=1}^{K_i} f(z_k \mid \boldsymbol{\alpha}_{ji}, Z_{k-1}) g_{ji,k} \end{aligned} \tag{3}$$

where $f(z_k \mid \boldsymbol{\alpha}_{ji}, Z_{k-1})$ is the conditional probability of the ji^{th} filter (defined by $\boldsymbol{\alpha}_{ji}$) and $P(\boldsymbol{\alpha}_{ji} \mid Z_{k-1})$ is the *a priori* probability among filters in the i^{th} bank that the j^{th} filter models the input generating environment. The bank level GN approximates this *a priori* probability in the gating weight $g_{ji,k}$. Assuming a Gaussian distribution on the measurement residuals of each filter, the probability distribution of the ji^{th} filter is given by

$$f(z_k \mid \boldsymbol{\alpha}_{ji}) = \frac{1}{\sqrt{2\pi W_{ji,k}}} e^{-\frac{r_{ji,k}^2}{2W_{ji,k}}} \tag{4}$$

where $r_{ji,k}$ is the prefit measurement residual and $W_{ji,k}$ is the innovations covariance of the ji^{th} filter at time t_k. The prefit measurement residual is given by

$$r_{ji,k} = z_k - h(\hat{\boldsymbol{x}}_{ji,k}^{(-)}; \boldsymbol{\alpha}_{ji}) \tag{5}$$

where $h(\hat{\boldsymbol{x}}_{ji,k}^{(-)}; \boldsymbol{\alpha}_{ji})$ is the measurement computed from the Kalman filter model evaluated at the current state estimate $\hat{\boldsymbol{x}}_{ji,k}^{(-)}$. The innovations covariance is a function of the state error covariance, $\boldsymbol{P}_{ji,k}^{(-)}$, the measurement mapping matrix, $\boldsymbol{H}_{ji,k}$, and the measurement noise covariance, $\sigma_{ji,k}^2$,

$$W_{ji,k} = \boldsymbol{H}_{ji,k} \boldsymbol{P}_{ji,k}^{(-)} \boldsymbol{H}_{ji,k}^T + \sigma_{ji,k}^2. \tag{6}$$

The prefit residual, the innovations covariance, and the measurement noise covariance are all scalar values in this implementation because the filter is processing scalar Doppler observables.

The likelihood function of the hierarchy may therefore be expressed as a function of the individual filter conditional probabilities and the top and bank level gating weights (*a priori* probabilities):

$$f(z_k) = \sum_{i=1}^{L} g_i \sum_{j=1}^{K_i} f(z_k \mid \boldsymbol{\alpha}_{ji}) g_{ji,k} \tag{7}$$

Note that explicit dependence on previous information, Z_{k-1}, will now be implicitly assumed to simplify the notation.

The gating weight *a priori* approximations are calculated by transforming the synaptic weights of the top and bank level GN neurons through the *softmax* operator [6, 7]. The gating weight of the i^{th} output cell of the top-level GN is defined using softmax as a function of the top level GN synaptic weights

$$g_i = \frac{e^{a_i}}{\sum_{j=1}^{L} e^{a_j}} \tag{8}$$

where a_i is the top level synaptic weight for each bank associated neuron [8, 9]. The gating weight assigned to the j^{th} filter within the i^{th} filter bank is similarly given by

$$g_{ji,k} = \frac{e^{a_{ji,k}}}{\sum_{n=1}^{K_i} e^{a_{ni,k}}}. \tag{9}$$

The softmax function serves to provide a differentiable activation function that preserves rank order and generalizes a winner-takes-all paradigm by exponentially separating gating weights. As an illustrative example, consider the synaptic and gating weights in Table 1. Rank order is preserved as the synaptic weights of the banks map into the same magnitude order in the gating weights. The winner-takes-all paradigm is observed when the relative gating weight magnitudes within the rank order are shifted to favor the bank with the largest synaptic weight. For example, the 50% increase in synaptic weight from bank 2 to bank 3 translates into a marked 271% increase in gating weight.

Bank	a_i	g_i
0	-2.0	0.003
1	0.0	0.0350
2	2.0	0.2583
3	3.0	0.7020
Sum	3.0	1.0

Table 1: Softmax Illustration

The gating weights are renormalized with each adaptation of the synaptic weights, which occurs when a measurement is processed by the Kalman filters, and have the following properties

$$0 < g_i < 1 \quad , \quad 0 < g_{ji} < 1 \tag{10}$$

$$\sum_{i=1}^{L} g_i = 1 \quad , \quad \sum_{j=1}^{K_i} g_{ji} = 1 \tag{11}$$

The gating weights may therefore be reasonably interpreted as the approximation to the conditional *a priori* probabilities [10] that each filter or bank of filters, depending on the level in the hierarchy, models the spacecraft dynamics and measurement environment correctly.

Adapting the synaptic weights (and hence modifying the gating weights) of the HME to maximize Eq. 7 is equivalent to adapting the synaptic weights to maximize the log-likelihood function

$$l = \ln f(z_k), \tag{12}$$

which proves easier to manipulate when evaluating a gradient with respect to the synaptic weights.

Substitution of the g's into Eq. 12 gives the log-likelihood in terms of the filter products and the gating weight vectors. The sensitivity of the log-likelihood to the synaptic weights is

$$\frac{\partial l}{\partial a_{i,k}} = h_{i,k} - g_{i,k}, \tag{13}$$

at the top level and

$$\frac{\partial l}{\partial a_{ji,k}} = h_{i,k}(h_{ji,k} - g_{ji,k}), \tag{14}$$

at the bank level. The h's are defined as

$$h_{i,k} = \frac{g_{i.k} \sum_{j=1}^{K_i} f(z_k \mid \boldsymbol{\alpha}_{ji}) g_{ji,k}}{\sum_{n=1}^{L} g_{n,k} \sum_{j=1}^{K_n} f(z_k \mid \boldsymbol{\alpha}_{jn}) g_{jn,k}} \tag{15}$$

for the top level and as

$$h_{ji,k} = \frac{f(z_k \mid \boldsymbol{\alpha}_{ji}) g_{ji,k}}{\sum_{j=1}^{K_i} f(z_k \mid \boldsymbol{\alpha}_{ji}) g_{ji,k}}. \tag{16}$$

for the ji^{th} filter on the bank level. Because of their dependence on filter products, the h's may be thought of as *a posteriori* probabilities of each filter or bank. It is through the *a posteriori* probabilities that the filter performances influence the GN learning process. This *a posteriori* probabilistic interpretation of filter outputs

has been used in other adaptive Kalman filter applications with varying degrees of success [11, 12].

The GN updates the synaptic weights $a_{i,k}$ and $a_{ji,k}$ after the experts process z_k by the gradient ascent procedure which seeks to maximize Eq. 12 [8]. The adaptation of the synaptic weights is

$$a_{i,k+1} = a_{i,k} + \eta \frac{\partial l}{\partial a_i} = a_{i,k} + \eta(h_{i,k} - g_{i,k}) \tag{17}$$

and

$$a_{ji,k+1} = a_{ji,k} + \eta \frac{\partial l}{\partial a_{ji,k}} = a_{ji,k} + \eta h_{i,k}(h_{ji,k} - g_{ji,k}) \tag{18}$$

where η is a learning rate parameter that is not generally the same for the top and bank level gating networks. The effective result of this learning scheme is that the gating weights track the *a posteriori* probabilities which are functions of filter residual and innovations covariance magnitudes. By adapting the synaptic weights to maximize Eq. 12, the best performing filters will generally be given a weight close to unity by the bank level GN. Similarly, the best performing bank will be given a weight close to unity by the top level GN; however, the presence of sufficiently poor performing filters in a bank containing the optimal filter may cause that bank to not receive the highest top-level gating weight. Note also that the bank level learning rule in Eq. 18 scales the update to the synaptic weight by the h value for that bank. Because of this scaling, internal learning progresses at a slower rate for filter banks that are assigned low top level gating weights.

EXPERIMENTAL RESULTS

The experimental application of the autonomous monitoring component is the MPF cruise segment beginning after TCM 2 on February 4, 1997 and ending April 17, 1997. Non-gravitational accelerations, SRP, and measurement noise have been identified as three of the most significant sources of uncertainty during MPF cruise [13]. Chaer and Bishop [2] were able to successfully identify changes in process noise statistics related to these error sources using an approximate dynamic model linear Kalman filter and simulated MPF data. Their results indicated that an HME implementation using high fidelity filters processing real data might be able to identify deterministic changes in these error sources. Therefore, the goal of the following experiments is to determine if a general HME configuration can be used to identify (1) non-gravitational accelerations in the form of unmodeled impulsive maneuvers, (2) SRP environment changes, and (3) unmodeled impulses and measurement noise statistics. The top level gating network must clearly assign weight to the correct macromode identification bank in the event of a spacecraft environment change. Also, the HME must be able to avoid false identification by giving the control filter bank the highest weight in the absence of environment changes. An initial learning

period is expected as all filters and banks are initially weighted equally in the HME. The learning rate for all GNs is set to 10 in the following experiments and all filters are configured with the operational parameter realization defined in Ref. [1] unless otherwise noted.

EXPERIMENT ONE: Impulsive Events

To identify unmodeled impulses, the HME must contain a specific bank of filters that will accommodate an unmodeled velocity change better than the operational filter in the control bank. Bank 0 in the HME configuration of Table 2 contains filters with impulsive maneuver filter states spaced evenly in time for the Doppler data span under consideration. These states are corrections to zero magnitude impulses applied at these times and should allow state space to accommodate unmodeled velocity changes by absorbing residual signal through the Kalman gain. The small MPF trajectory maneuver, approximately 0.7 mm/sec, on March 25, 1997 will be used as the impulse to be identified in this experiment. This maneuver magnitude is just above the 0.18 mm/sec noise level modeled in the operational filter Doppler measurements [1] and has been omitted from the model of all filters.

The filters in bank 1 use the MPF navigation team SRP solutions 2 and 4 as detectors for SRP change and the operational Doppler noise value of 0.01 Hz (0.18 mm/sec) [1]. The optimal SRP model is indicated as "Ely" and is a model tuned by the genetic algorithm in the adaptive component of navigation architecture [5]. Filters in bank 2 model changes in the Doppler measurement noise but are otherwise optimal. Filter (0,3) is the operational filter and is optimal other than the absence of the March 25 impulse.

Filter Number	Impulse	SRP Model	R
0,0	Feb.4 [0,0,0]	Ely	0.01
1,0	Feb.22 [0,0,0]	Ely	0.01
2,0	Mar.12 [0,0,0]	Ely	0.01
3,0	Mar.30 [0,0,0]	Ely	0.01
0,1	—	MPF 2	0.01
1,1	—	MPF 4	0.01
0,2	—	Ely	0.003
1,2	—	Ely	0.03
2,2	—	Ely	0.09
0,3	—	Ely	0.01

Table 2: Impulsive Event Experiment HME Configuration

The top level gating weights from this experiment are provided in Figure 3 where the times of the unmodeled impulse and the bank 0 test impulse states are indicated by vertical lines. The gating weights are equally distributed between the impulse and

control banks for the first 10 days of data. This result is to be expected as three of the filters in bank 0 contain the same dynamic model as the operational filter over the first quarter of the data span. However, when the February 22 test impulse in filter (1,0) comes and goes without an impulse occurring, weight is reassigned from the impulse test bank to the control bank. The autonomous monitor passes an important test at this point by avoiding a false detection of a macromode change. A similar effect is observed on March 12 to a smaller extent. The unmodeled impulse occurs in the middle of the data pass on day 49 but does not manifest a significant enough change in the data for identification until the next pass when the impulse bank is correctly assigned the majority of the gating weight. A navigator examining this gating history would be justified in adapting the operational filter to accommodate an impulse between March 12 and March 27. The bank level gating weights in Figure 3 reflect a similar behavior as filter (2,0) is assigned the most weight in bank 0 on the pass after the unmodeled impulse. The SRP alternative models in bank 1 are weighted relatively the same within their bank. The small noise filter (0,2) is selected throughout the data span within bank 2. However, recall from Eq. 18 that bank level learning is scaled by the top level *a posteriori* weight, $h_{i,k}$, associated with each bank. Since the gradient learning rule in Eq. 17 indicates that the gating weights track the *a posteriori* weights and the noise bank gating weight goes to zero in Figure 3, it can be concluded that relatively no internal learning occurs in bank 2 because of its poor collective top level performance.

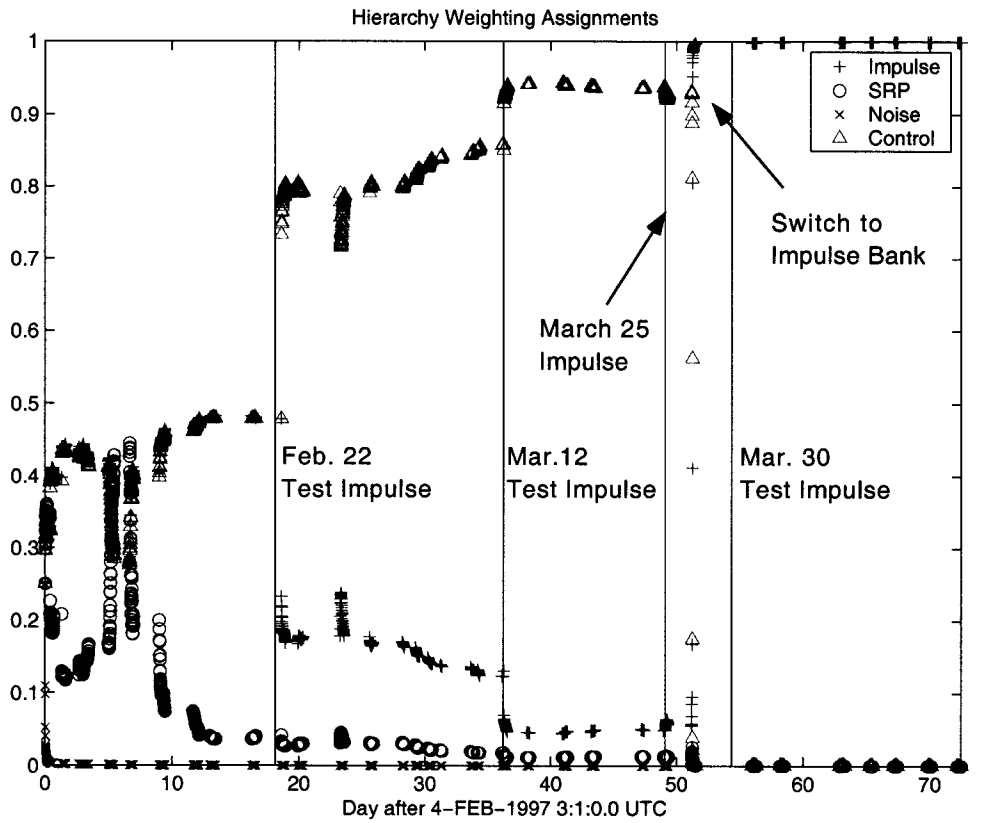

Figure 3: Impulse Event Top Level Weights

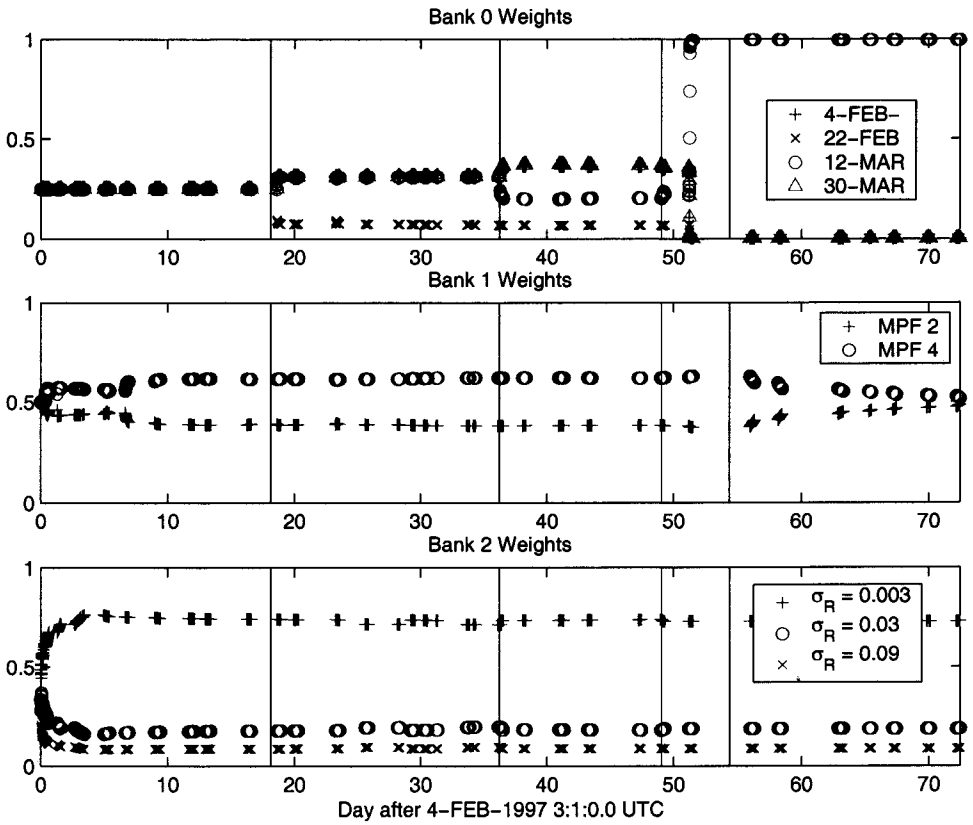

Figure 4: Impulse Event Bank Level Weights

EXPERIMENT TWO: SRP Model Changes

In order to investigate the ability of the autonomous component to identify a continuous change in the SRP model, the HME was configured as detailed in Table 3. The basic structure is the same as in the impulsive event detection with the exception that all of the filters not involved with SRP change identification utilize the MPF navigation team SRP model 4. The assumption of this scenario is that MPF SRP model 4 is incorrectly believed to be the optimal model and is therefore used in the operational filter and most other filters. The two filters in the SRP change bank 1 use the "Ely" genetic algorithm tuned SRP model and the MPF navigation team model 2 to provide alternative SRP dynamics to the operational model. The Doppler noise values are configured in the same manner as before with bank 2 filters modeling changes in the measurement noise. To examine the degree to which the top level GN can switch between macromodes, the small March 25 impulse is again omitted from all filters.

Filter Number	Impulse	SRP Model	R
0,0	Feb.4 [0,0,0]	MPF 4	0.01
1,0	Feb.22 [0,0,0]	MPF 4	0.01
2,0	Mar.12 [0,0,0]	MPF 4	0.01
3,0	Mar.30 [0,0,0]	MPF 4	0.01
0,1	—	MPF 2	0.01
1,1	—	Ely	0.01
0,2	—	MPF 4	0.003
1,2	—	MPF 4	0.03
2,2	—	MPF 4	0.09
0,3	—	MPF 4	0.01

Table 3: SRP Model Change Experiment HME Configuration

After the initial 10 day learning period, the top level gating weights are definitively assigned to the SRP model change bank as illustrated in Figure 5. Examining the bank level weights in Figure 6, it can be seen that the optimal "Ely" SRP model in expert (0,1) receives the majority of weight from the beginning of the data span but also gains a definitive weighting majority after day ten. The top level SRP bank maintains the maximum weighting through the test impulses on February 12 and March 22 but the impulse bank is selected beginning on the data pass after the unmodeled impulse. Therefore, in this experiment the autonomous monitor is able to distinguish between a continuous and a discrete environment change within the same data set.

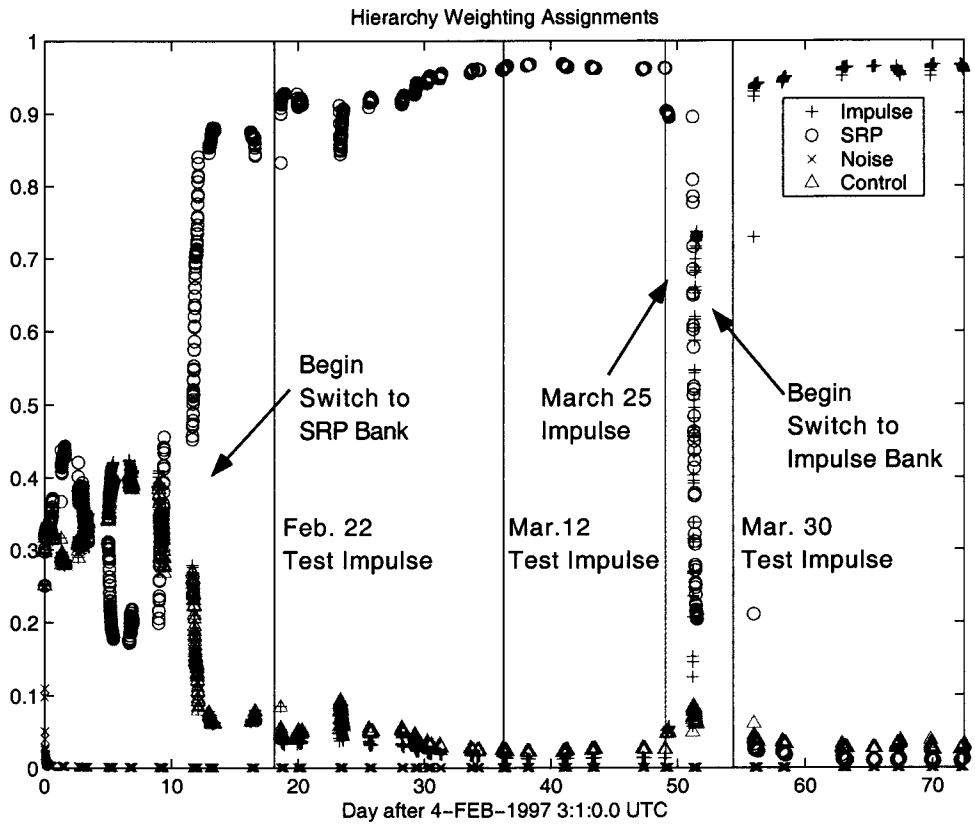

Figure 5: SRP Change Top Level Weights

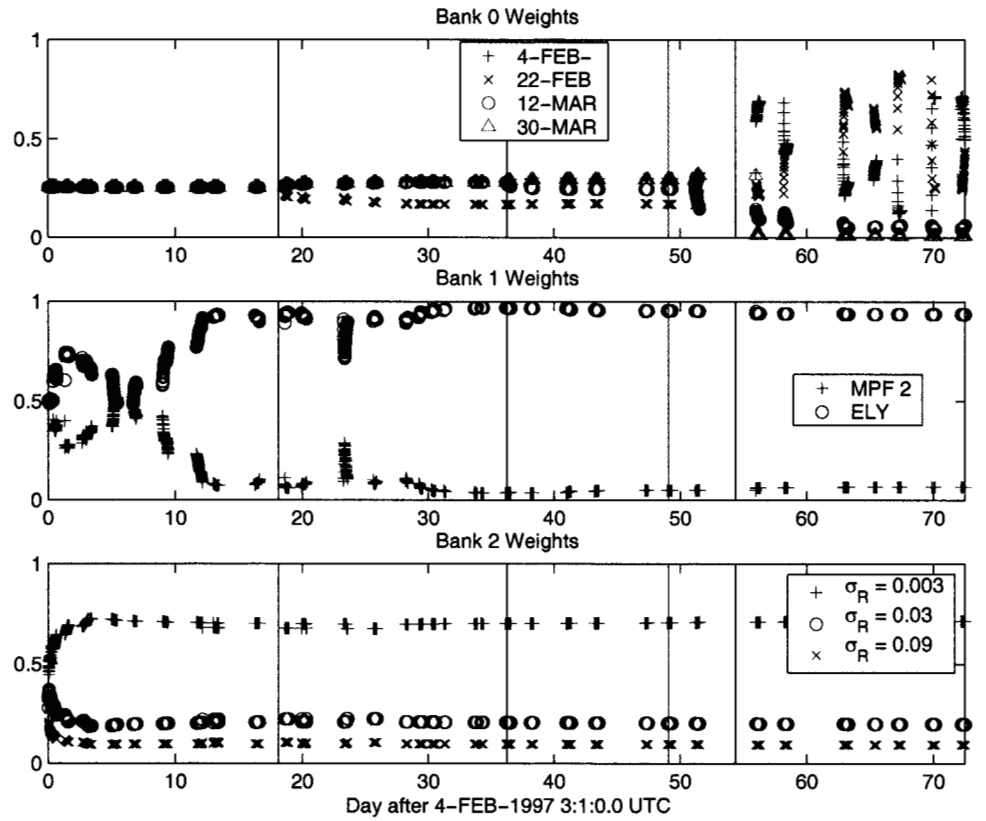

Figure 6: SRP Change Bank Level Weights

EXPERIMENT THREE: Noise Change and Impulsive Event Detection

The impulsive event detection performed in the first experiment is now conducted with the addition of the Gaussian noise signal plotted in Figure 7. The addition of this noise from day 5 to day 15 gives the measurements an equivalent total noise value of 0.04 Hz for this time period. The HME is configured exactly the same as in the impulse detection experiment as illustrated in Table 4.

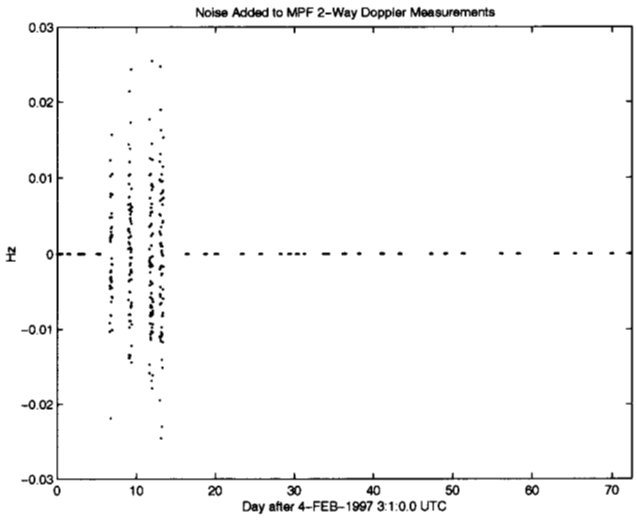

Figure 7: Noise Added to MPF Doppler Data

Filter Number	Impulse	SRP Model	R
0,0	Feb.4 [0,0,0]	Ely	0.01
1,0	Feb.22 [0,0,0]	Ely	0.01
2,0	Mar.12 [0,0,0]	Ely	0.01
3,0	Mar.30 [0,0,0]	Ely	0.01
0,1	—	MPF 2	0.01
1,1	—	MPF 4	0.01
0,2	—	Ely	0.003
1,2	—	Ely	0.03
2,2	—	Ely	0.09
0,3	—	Ely	0.01

Table 4: Impulse and Noise Experiment HME Configuration

The top level and bank level weights for the noise and impulse identification experiment are provided in Figures 8 and 9. The top level weights exhibit a brief initial learning period before the noise bank moves to almost unity weighting very close to the beginning of the increased noise segment. On the bank level, filter (1,2)

with a noise statistic of 0.03 Hz moves to unity weighting over the 0.003 and 0.09 Hz filters. The increased noise segment ends on day 15 but the noise bank maintains unity top level weight until the learning period starting near day 30. It is not clear whether the extended weighting of the noise bank is a property of gating network or the EKFs or both. However, during the learning period the control and impulse banks are weighted almost equally. Once the March 12 test impulse in filter (2,0) passes near day 35, the control bank is definitively selected as the best performer. When the unmodeled impulse occurs on March 25, a second learning period occurs where the noise bank is given the highest weight for two data passes. After March 30, the impulsive filter is given unity weighting. The soft decision capability of the multi-level HME is demonstrated in this switch because the March 12 test impulse filter is selected after the unmodeled impulse, only to give way to the March 30 test impulse filter on the bank level, but the top level decision is nonetheless correct.

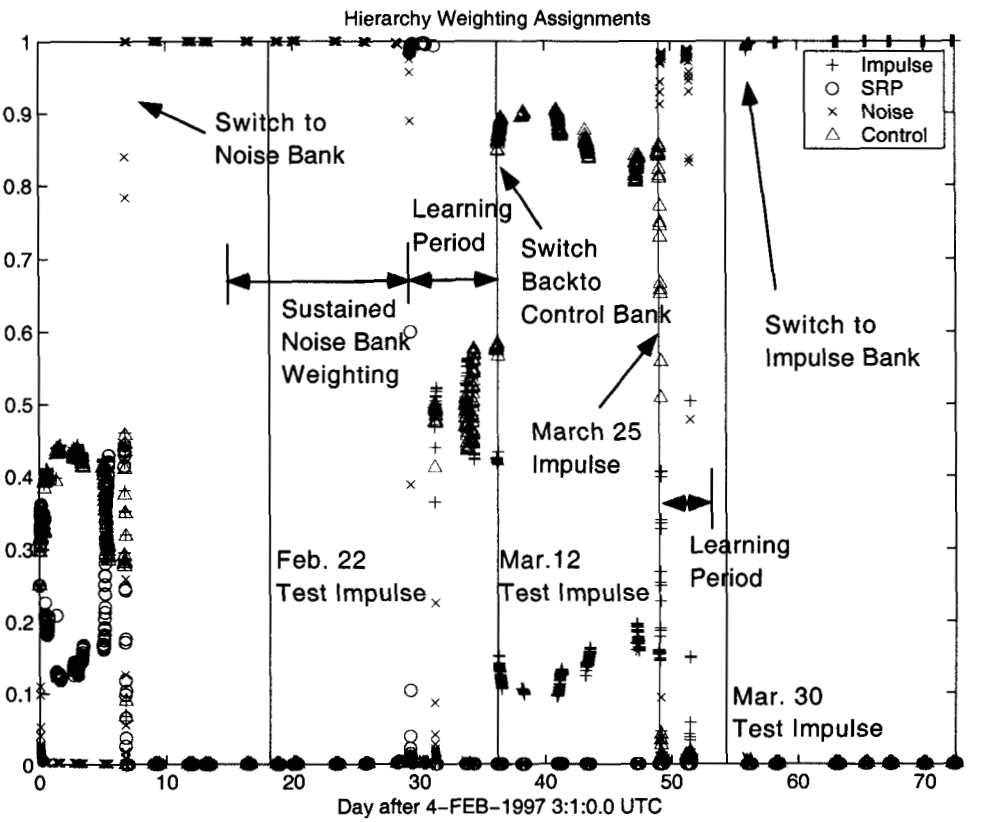

Figure 8: Noise and Impulse Top Level Weights

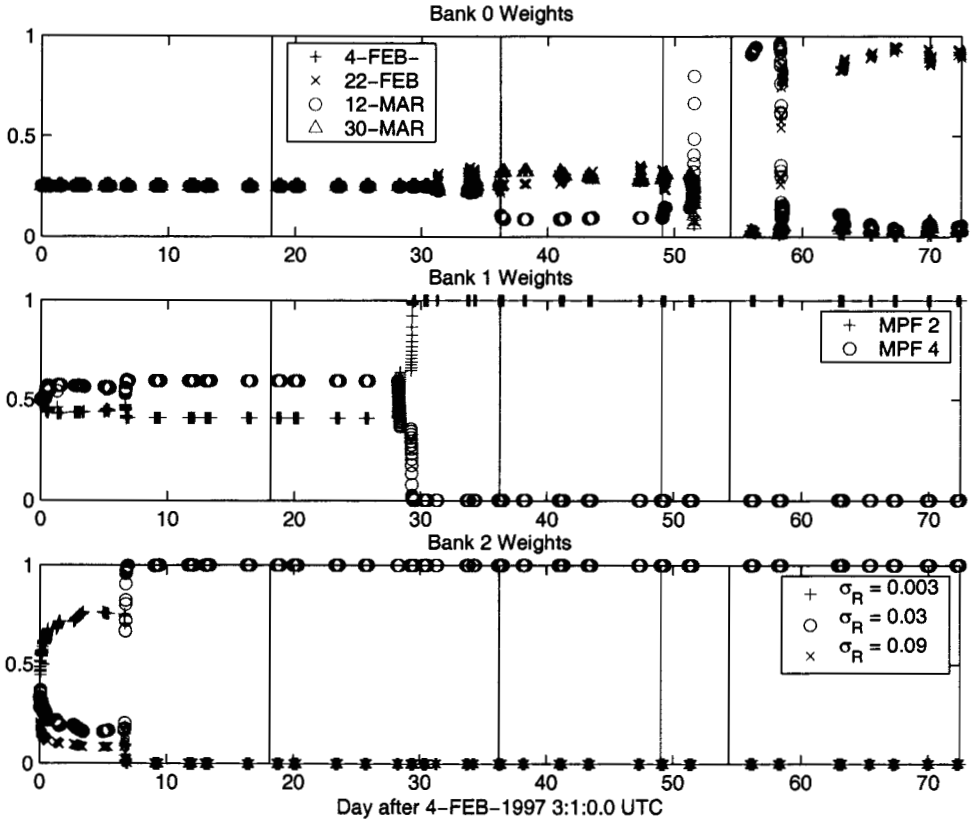

Figure 9: Noise and Impulse Bank Level Weights

SUMMARY AND CONCLUSIONS

The autonomous monitoring algorithm successfully detected and identified unmodeled macromode environment changes in the MPF tracking environment within one or two data passes. Changes were identified correctly in all three experiments and false identifications in the absence of environment changes were not significantly present. Learning periods were observed at the beginning of each experiment when no filter bank was clearly selected or the incorrect filter bank was temporarily selected by the top level gating network. The multi-mode MPF environment changes in the third experiment created two distinct learning periods within the data span after each mode change and incorrectly maintained noise macromode selection for approximately 15 days after the increased noise period elapsed. Some learning is required for the top level gating network to make robust decisions, otherwise a single measurement could cause an instantaneous reassignment of gating weight and lead to false detection or incorrect identification of environment changes. The learning periods at the beginning of the data span might be reduced if the gating networks were initialized with synaptic weights favoring the operational filter in the control bank instead of equal initial weighting as in these experiments. Such a biased initial weighting scheme has intuitive merit since the operational filter is believed to be the best model of the spacecraft environment when the first measurement is processed. However, since the adaptation of the synaptic weights is sensitive to their magnitude during application of the gradient learning rule, care must be taken to not overlook changes that occur shortly after the data span begins. As always, the balance between response time and robust decision making must be observed.

Further experimentation is required to determine the effectiveness of the impulse identification configuration used in bank 0. In particular, the temporal spacing of test impulse filter states required to adequately determine the approximate time of unmodeled discrete velocity changes in an operational setting has yet to be determined. It is possible that the adaptation of the operational filter to account for unmodeled impulses might require the genetic algorithm to include event time as one of its parameters. The work by Gholson [11] with a single linear Kalman filter modeling multiple unknown accelerations inspires the creation of a queue of impulse detection EKFs in bank 0. In the queued approach, a manageable number of filters populate bank 0 with a reasonable temporal test impulse spacing. If bank 0 does not accrue top level weight within a certain time interval after the first test impulse in the queue has occurred, then that filter is removed from the bank and a new filter with an upcoming test impulse replaces it at the end of the queue. The multi-level structure of the HME would be useful in this approach since bank level probabilities would be reassigned equally when the queue changes, but the top level weight assigned to the impulse detection bank would remain the same. This approach is appealing as it limits the number of active filters required for small temporal spacing of test impulses in bank 0.

Several operational questions must be addressed before the autonomous mon-

7. Bridle, J.S., "Training Stochastic Model Recognition Algorithms as Networks can Lead to Maximum Mutual Information Estimation of Parameters," *Advances in Neural Information Processing Systems 2*, ed. D.S. Touretzky, pp. 211-217, Morgan Kaufmann, 1990.

8. Haykin, S., *Neural Networks a Comprehensive Foundation*, Prentice-Hall Inc., 2nd ed., Upper Saddle River, New Jersey, 1999.

9. Crain, T.P. and Bishop, R.H., "Unmodeled Impulse Detection and Identification during Mars Pathfinder Cruise," Center for Space Research Technical Memorandum, CSR-TM-00-01, March 2000.

10. Haykin, S., *Neural Networks a Comprehensive Foundation*, Macmillan Publishing Company, 1st ed., New York, 1994.

11. Gholson, N.H. and Moose, R.L., "Maneuvering Target Tracking Using Adaptive State Estimation," IEEE Transactions on Aerospace and Electronic Systems, Vol. AES-13, No. 3, May 1977, pp. 310-317.

12. Burkhart, P.D. and Bishop, R.H., "Adaptive Orbit Determination for Interplanetary Spacecraft," *Journal of Guidance, Control, and Dynamics*, Vol. 19, No. 3, AIAA, Washington DC, May-June 1996, pp.693-701.

13. Burkhart, P.D., Bishop, R.H., and Estefan, J.A., "Covariance Analysis of Mars Pathfinder Interplanetary Cruise," *Advances in the Astronautical Sciences*, edited by R.J. Proulx, J.J.F. Liu, P.K. Siedelmann, and S. Alfano, Vol. 89, Pt. 1, Univelt, San Diego, CA, 1995, pp.267-285.

itoring and the offline adaptive components can be fully integrated. Given that periods of learning can occur even with acceptable top level decision making, two distinct problems arise. First, a method of distinguishing ongoing learning from environment changes beyond the scope of those modeled in the HME configuration of filters should be developed. It is possible that the HME can diverge if an change occurs in which no filter or bank correctly models the spacecraft environment. Large residuals in all filters can cause a numerical underflow problem in the *a posteriori* probability calculations. A time limit or data quantity limit on learning along with minimum residual magnitude monitoring should adequately address this issue. Second, a method of determining when adaptation should take place should be developed. The operational filter should be adapted by the genetic algorithm only after the top level weights have maintained a threshold value for a specified number of data points. This would avoid adaptation in response to transient or learning period weighting regimes.

ACKNOWLEDGMENTS

The work described in this paper was partially supported Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration and by a grant from the Center for Space Research at the University of Texas at Austin.

REFERENCES

1. Vaughan, R. M., Kallemeyn, P. H., Spencer, D. A., and Braun R. D., "Navigation Flight Operations for Mars Pathfinder," AAS 98-145, February 1998.

2. Chaer, W.S., Bishop, R. H., and Ghosh J., "Hierarchical Adaptive Kalman Filtering for Interplanetary Orbit Determination," IEEE Transactions on Aerospace and Electronic Systems,Vol. 34, No. 3, 1998, pp. 1 -14.

3. Crain, T.P. and Bishop, R.H., "The Mixture-of-Experts Gating Network: Integration into the ARTSN Extended Kalman Filter," Center for Space Research Technical Memorandum, CSR-TM-99-01, March 1999.

4. Chaer, W.S. and Bishop, R. H., "Adaptive Kalman Filtering with Genetic Algorithms," *Advances in the Astronautical Sciences*, edited by R.J. Proulx, J.J.F. Liu, P.K. Siedelmann, and S. Alfano, Vol. 89, Pt. 1, Univelt, San Diego, CA, 1995, pp.141-156.

5. Ely, T.A., Bishop, R.H., Crain, T.P., "Adaptive Interplanetary Navigation using Genetic Algorithms", To appear in *The Journal of Astronautical Sciences*, 2000.

6. Bridle, J.S., "Probabilistic Interpretation of Feedforward Classification NEtwork Outputs, with Relationship to Special Pattern Recognition," *Neurocomputing: Algorithms, Architectures, and Applications*, eds. F. Fougleman-Soulie and J. Herault, pp. 227-236, Springer-Verlag, 1990.